

ESTIMATING THE PHASE CONGRUENCY OF LOCALISED FREQUENCIES

Peter J. Myerscough, Mark S. Nixon

School of Electronics and Computer Science
University of Southampton, UK, SO17 1BS

ABSTRACT

Phase congruency is a new method for detecting features in images. One of its significant strengths is its invariance to lighting variation within an image, as well as being able to detect a wide range of interesting features. We present a method for estimating the phase congruency of localised frequencies that cannot be measured separately by Gabor filters. We show that by measuring the ratio of the standard deviation to the mean energy between different phase shifted Gabor filters that we are able to estimate whether the localised frequencies are phase congruent. We then show example results from applying this estimation procedure to a set of images. We also show improvements when compared to another phase congruency detector. We conclude that the concept of estimating the phase congruency of localised features is possible, but more work is needed to mature the technique to a robust feature detector.

1. INTRODUCTION

Phase congruency can be used to detect features within images. It can be explained by examining a simple square wave as in fig. 1(a) which shows a series of sine-waves, that when summed together form a square wave. What is of interest is that at the rising and falling edges of the square wave the sine-waves all have the same phase, or are phase congruent. This behaviour is also seen in the peak and trough of the triangle wave in fig. 1(b) and at the centre of the impulse function in fig. 1(c).

The benefits of using a measure of phase congruency to find features is its illumination invariance and good localisation. There has been some work in implementing methods that measure phase congruency in images[1, 2] and image sequences[3, 4]. From our own work we have found that 'faster' moving features can cause a band of feature responses around the true features using the current phase congruency techniques. These techniques measure the phase congruency between filter responses, but have no knowledge of the phase congruency of the frequency components that are within the bounds of the filters they use. We have developed a new technique for estimating the phase




Fig. 1. Example waveforms showing high phase congruency highlighted by circles.

congruency of the frequency components within the bounds of a filter which does not exhibit this banding behaviour.

Our new technique relies on localised estimates of the phase congruency of localised frequencies. This differs from Gabor filtering which gives localised phase information of a group of frequencies. To estimate this type of phase congruency we apply linear and non-linear phase shifts to a Gabor filter and examine the ratio of the standard deviation to mean energy of the responses. This allows the successful separation of frequencies with high and low degrees of phase congruency, and gives clear indications as to the presence of an image feature. The technique is then tested on a set of images with some conclusions and ideas for future work.

2. ESTIMATING THE PHASE CONGRUENCY OF LOCALISED FREQUENCIES

Using Gabor filters it is possible to measure the localised phase congruency of groups of frequencies. There is no better method to gain spatial localisation and frequency separation than to use Gabor filters. Therefore to develop a technique to estimate the phase congruency of frequencies we have used a framework that mimics the grouping and summing of localised frequencies when calculating a Gabor filter response. Gabor filter responses can be written as in eqn. 1.

$$\begin{aligned}
G_c &= \sum_n^m a_n R_n \cos(\theta_n) \\
G_s &= \sum_n^m b_n R_n \sin(\theta_n)
\end{aligned} \quad (1)$$

where a_n and b_n are the filter coefficients, R_n is a localised frequency's amplitude, and θ_n is the localised frequency's phase. G_c and G_s are the cosine and sine Gabor filter responses. If two elements of the sum are π radians out of phase, they will cancel each other out. Unless all the elements are perfectly in phase, the sum of the absolute energy of the elements will always be greater than the sum of the elements. Mathematically to solve the above equation with respect to finding all the θ_n 's, we would need to repeat the Gabor filtering process many times with different coefficients. This is impractical and would also resolve the amplitude of each individual localised frequency, information we are not interested in. Instead we note the relationship between the phase congruency of the localised frequencies and the sum of energy. It is known that phase congruent frequencies when summed give greater responses than random phased frequencies that will negatively superpose. If different phase shifts were added to the Gabor filter's coefficients then the sum of energy will change. Perfectly congruent frequencies would be expected to reduce in the amplitude of response, and random phase frequencies remain of approximately the same amplitude, as illustrated in eqn. 2.

$$\begin{aligned}
G_c &= \sum_n^m a_n \cos(\psi_n) R_n \cos(\theta_n) \\
G_s &= \sum_n^m b_n \sin(\psi_n) R_n \sin(\theta_n)
\end{aligned} \quad (2)$$

where ψ_n represents the different phase angles of the Gabor filter's coefficients. Since the exact distribution of energy amongst the localised frequencies will be unknown, a single phase shift may not significantly affect the response to the filters. It is necessary then to repeat the filtering L times, forcing different amounts of phase shift to occur at different frequencies. To extract a measure of the effect of the phase shifting, we calculate the ratio of the standard deviation to the mean energy of the phase shifted responses. Since the energy of the localised frequencies will have been mixed between cosine and sine-based Gabor filter responses, we use the energy from the filters, E .

$$E_l = \sqrt{G_{c_l}^2 + G_{s_l}^2} \quad (3)$$

$$\rho_E = \frac{1}{\bar{E}} \sqrt{\frac{1}{L} \sum_l^L (E_l - \bar{E})^2} \quad (4)$$

where E_l is the energy from the l th phase shifted filter. \bar{E} is the mean energy of the L different responses, and ρ_E is the ratio of the standard deviation of the different phase shifted filter responses and the mean energy.

To test whether this ratio describes the phase congruency of localised frequencies we generated sets of synthetic localised frequencies. Each set has a random amplitude spectra and phase angle, with increasing measures of zero mean Gaussian noise perturbing the phase of the localised frequencies. If the ratio, ρ_E , is to be a measure of phase congruency it should be correlated to the change in noise levels.

Before testing, the phase shifts applied to the Gabor filters need to be defined. If we look at randomly shifting each frequency element in the Gabor filter the spatial extent of the filter would extend significantly. Therefore the phase shifts applied to the Gabor filter should affect the spatial extent as little as possible. There should also be different amounts of shift applied to each frequency. The first set of tests used linear phase shifts ($\psi_n = 2\pi n/m$), but additional phase shift patterns were investigated. The final design for the phase shifts is non-linear and based on the square root of the cosine and sine functions. This series of phase shifts has a good balance between spreading the Gabor filter spatially and providing useful changes in phase across the length of the Gabor filter. It is also normalised to avoid falsely increasing or decreasing the detected energy.

$$\begin{aligned}
\psi_n &= 2\pi n/m + 2\pi l^2/L^2 \\
G_{c_l} &= \sum_n^m a_n \frac{|\sqrt{\cos(\psi_n)}|}{\sqrt{|\sqrt{\cos(\psi_n)}|^2 + |\sqrt{\sin(\psi_n)}|^2}} R_n \cos(\theta_n) \\
G_{s_l} &= \sum_n^m b_n \frac{|\sqrt{\sin(\psi_n)}|}{\sqrt{|\sqrt{\cos(\psi_n)}|^2 + |\sqrt{\sin(\psi_n)}|^2}} R_n \sin(\theta_n)
\end{aligned} \quad (5)$$

2.1. APPLYING PHASE CONGRUENCY OF LOCALISED FREQUENCY

Having developed a localised phase congruency detector it is possible to detect features in images based upon the ratio of standard deviation to mean energy of a set of filter responses. Results from just using the ratio from the localised frequencies also included points in the image where no features or image structure are present. This is accounted for by the low energy or zero energy at those points and therefore the small standard deviation calculated at these points. A basic approach to this problem is to divide the mean energy from the filters by the ratio and then normalise the result based on the greatest value.

$$n(x, y) = \frac{\bar{E}(x, y)}{\rho_E(x, y)}$$



Fig. 2. These figures show three original images and their feature detected counterparts. The scales to the right of each image show the feature strength.

$$\begin{aligned}
 d(x, y) &= \max(n(x, y)) \\
 pc(x, y) &= \frac{n(x, y)}{d(x, y)} \quad (6)
 \end{aligned}$$

where $\max(n)$ is the maximum value over all n in the image.

3. RESULTS

3.1. LOCALISED FREQUENCY PHASE CONGRUENCY

We tested the ratio as a method to estimate the phase congruency of the localised filters on a set of 10,000 synthetic frequency sets. The results are plotted in fig. 3 and show eqn. 5 has a greater ability to distinguish high and low phase congruency than eqn. 2.

Eqn. 6 was used to generate feature maps for a set of images. Example results are shown in figures 2. In fig. 2(d) the woman's outline is well detected with a positive attempt at detecting the edge of the hair and the features within the bunch of flowers. In fig. 2(e) the outline of each houses' roof and the top half of the trunk of the tree is quite visible as




Fig. 3. This plots show the ratio, ρ , plotted against localised frequency sets with increasing amounts of phase perturbation. The solid line shows the results of applying eqn. 2 and the dashed line that of eqn. 5. The central line, in each case, joins the median values with the 25% quartile and the 75% quartile projected above and below. Tests were run on with random mean phase angles and different phase perturbations a series of 10,000 times. Results for eqn. 5 show a strong correlation between the ratio, ρ , and the variation in phase congruency produced by the phase perturbation.

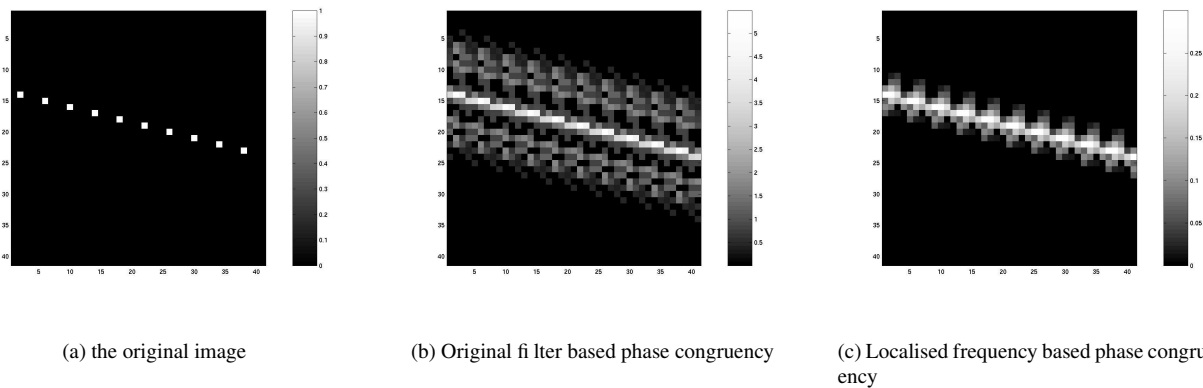


Fig. 4. These figures show two feature detected images of a dot 'moving' in a 1D signal through time or of a line of dots (depending on your labeling of the axes). The scale to the right shows the feature strength to grayscale mapping.

are the small branches protruding from the tree. In fig. 2(f) the outline of the legs, body and head are clearly extracted. The features about the sides of the body are almost lost due to the similar grayscales of the clothing and the background. These results certainly illustrate that the localised frequency phase congruency can be used for feature detection. The results have been non-maximally suppressed, and peaks have been capped to allow for easier viewing.

3.2. COMPARISON TO OTHER PHASE CONGRUENCY TECHNIQUES

One of the main motivations behind developing this technique for measuring phase congruency is the results from feature detecting a plain image with a line of dots traversing it, for example in fig. 4. The technique by Kovessi[1, 2] has some odd behaviour, as does some of my work on moving features[4].

In both images there is a central ridge with peaks on the exact position of the dots in the original image. The significant difference is that fig. 4(b) has two additional lines of dots/peaks along each side of the central line which give false feature points even after non-maximally suppressed.

4. CONCLUSIONS

We have developed a method that to estimate the phase congruency of localised frequencies. We have used this method to detect features via phase congruency in a number of images, in particular we have managed to remove an odd effect when detecting a line of dots. This should help to detect features that move through an image sequence, but are 'disconnected'. As a feature detector the method needs some improvements to remove the dependency on the amount of

energy a filter detects. It may also be useful to build a probabilistic framework to interpret the standard deviation data. This work will hopefully also benefit work with temporal phase congruency where moving features are extracted, and often disconnected.

5. ACKNOWLEDGEMENTS

This research has been funded by the EPSRC, with additional support from European Research Office of the US Army, Contract No.N68171-01-C-9002. Thanks also go to Ohio State university for some of the images shown.

6. REFERENCES

- [1] P Kovessi. Image features from phase congruency. *Videre : Journal of Computer Vision Research*, 1(3):1–27, 1999.
- [2] P. Kovessi. Phase congruency detects corners and edges. In *DICTA*, Sydney, December 2003.
- [3] M. Mulet-Parada and J.A. Noble. 2d+t acoustic boundary detection in echocardiography. *Medical Image Analysis*, 4:21–30, 2000.
- [4] P.J. Myerscough and M.S. Nixon. Temporal phase congruency. In *SSIAI*, 2004.